

REMARKS

The Examiner has requested cancellation of the withdrawn claims; has objected to the Specification; has objected to the Figures; has rejected Claim 28 under 35 USC 112; has rejected Claims 1-6, 11, 12, 15-20, 25-26, 28 and 30 under 35 USC 102 as anticipated by Li; has rejected Claims 7, 8, 13, 14, 21, 22, and 27 under 35 USC 103 as unpatentable over Li in view of Ganz; and, has rejected Claims 1-8, 11-22, 25-28, and 30 under 35 USC 102 as anticipated by Chirashnya. Applicants respond to the Examiner's requests herein and respectfully assert that the claims are allowable over the cited art.

With regard to the withdrawn claims, Applicants herein cancel those claims without prejudice to future prosecution of the claims. With regard to the objections to the Specification, Applicants have amended the Specification herein to provide a new Title of the Invention and to submit the data for a referenced patent application, which has not been available at the time of filing. With regard to the Figures, Applicants submit formal drawings herewith. In response to the rejection of Claim 28 as lacking antecedent basis, Applicants amend the preamble language of Claim 28, as well as Claim 30, to recite "a method" instead of "method steps."

Claims 1-6, 11, 12, 15-20, 25-26, 28 and 30 have been rejected under 35 USC 102 as anticipated by Li. The Li patent id
AUS920000829-US1

directed to a method for providing admission control (AC) for service requests based on bandwidth. Li defines "...an admit limit (AL) representing a bandwidth utilization capacity allocated to a service offered on the weakest link of the network". When a request for the service is received, the required bandwidth for the request is compared to the bandwidth utilization capacity for that service. If the required bandwidth exceeds the bandwidth utilization capacity, the request is rejected.

Applicants respectfully assert that the Li patent does not teach or suggest the invention as claimed. The present invention, as taught and claimed, provides a method, system, and program storage device for performing method steps for identifying slow links in a distributed network. Under the present invention, a first step is defining an original link speed factor for each of the links in the distributed network. What Li does is define a bandwidth utilization capacity for a particular service based on the service using the weakest link in Li's network. Li does not define an original link speed factor for each link. The term "link speed" refers to a specific, measurable value, as set forth in the Specification on page 2, and does not encompass all possible values assignable to a network. Applicants respectfully assert that bandwidth capacity is not the same as link speed.

Applicants further note that Li does not perform at least one runtime measurement of at least one runtime link speed indicator for each of a plurality of links. Li uses its original bandwidth capacity for a service throughout its operation. While Li can "deduct" capacity from that value, based on acceptance of requests, Li does not actually measure capacity during runtime, let alone measure a runtime link speed indicator. Li compares all incoming requests, routed from the edge routers to the QoS manager, to the originally determined bandwidth utilization capacity for the requested service. The bandwidth capacity is determined based on an assumption that a particular link is the weakest link, and will always be the weakest link in the network. Li does not teach or suggest dynamically determining if a link is weak.

Applicants further contend that Li does not teach or suggest the step of calculating a runtime link speed factor based on runtime measurements. As noted above, Li relies on the predefined bandwidth capacity and does not use measurements to dynamically assess runtime link speed. Accordingly, it cannot be concluded that Li calculates a runtime link speed factor based on measurements if Li has not performed any measurements related to link speed.

Finally, with regard to the claim feature of comparing the original link speed factor to the runtime link speed factor, Applicants reiterate that Li does not define an original link

speed factor, does not measure runtime link speed and does not calculate a runtime link speed factor. Clearly, therefore, Li cannot then compare values which Li has not defined, measured, or calculated.

For a reference to anticipate claim language under 35 USC 102, that reference must teach each and every claim feature. Since the Li patent does not teach the claimed steps or means for defining an original link speed factor for each link, performing at least one runtime measurement for each link, calculating a runtime link speed factor for each link, and comparing the calculated runtime link speed factor to the original link speed factor, it cannot be concluded that Li anticipates the invention as claimed in independent Claims 1, 17 and 28, or the claims which depend therefrom and add limitations thereto (Claims 2-8, 18-22), or those claims which recite parallel limitations (Claim 15).

With regard to the language of independent Claims 11, 25 and 30, and the claims which depend therefrom (Claims 12-16 and 26-27) and those claims which also recite application-based response to detected slow links (Claims 4-6, 8, and 20), Applicants again note that the Li patent does not teach the claimed step or means for detecting at least one slow link in the distributed network. Applicants rely on the arguments set forth above with respect to that claim feature. Further, Applicants contend that the Li patent neither teaches nor suggests the step

for each detected slow link of determining what specific applications require access to the detected slow link and adjusting application usage of the detected slow link by the specific applications. The Examiner has concluded that the Li passage, from the Abstract, which mentions "dynamic bandwidth adjustment" anticipates "dynamically adjusting application usage". Applicants respectfully disagree. The only dynamic bandwidth adjustment performed by Li is deducting required bandwidth for a request from the bandwidth utilization capacity when a request is accepted. Li does not dynamically adjust bandwidth based on measured or detected bandwidth, and clearly does not dynamically adjust based on measured link speed. Moreover, what Li teaches is that requests are either accepted or rejected based on the bandwidth required for the request as compared to the bandwidth utilization capacity which was predefined based on a predetermined weakest link in the network. Li does not teach or suggest adjusting application usage of links. Applicants reiterate that anticipation under 35 USC 102 can only be maintained if the reference teaches each and every claim feature. Li does not teach or suggest adjusting application usage, either by a system administrator or the application itself, in response to dynamic detection of slow links. Accordingly, Applicants conclude that the Li patent does not anticipate the language of Claims 11-16, 25-27, 30, or those

claims which also recite application-based response to detected slow links (Claims 4-6, 8, and 20).

Claims 7, 8, 13, 14, 21, 22, and 27 have been rejected under 35 USC 103 as unpatentable over Li in view of Ganz. Applicants rely on the arguments set forth above with regard to the teachings of the Li patent. Further, Applicants assert that the Ganz patent does not provide those teachings which are missing from the Li patent. Ganz is cited for its teachings related to an administrator identifying slow links and altering application usage of slow links. Applicants first note that the claim language does not recite that a system administrator identifies slow links. Rather, the claim language calls for the system administrator to be notified of designated slow links and to perform altering of application usage of the designated slow links. Moreover, the Ganz patent does not provide that a system administrator alter application usage of slow links based on dynamic detection of slow links based on link speed measurements. Ganz, like Li, looks to bandwidth capacity, which is not the same as or suggestive of link speed.

Claims 1-8, 11-22, 25-28, and 30 have been rejected under 35 USC 102 as anticipated by Chirashnya. The Chirashnya patent is directed to a system and method to detect faulty switch adapters. Chirashnya has multiple nodes transmit packets through a switch adapter which is to be tested and then detects, at the packet destination, whether a bad packet has been received. If a bad

packet is detected, the source (i.e., the faulty switch adapter) is identified. Alternatively, the packets arriving at the destination are counted, and a faulty switch adapter is identified if fewer packets arrived than were sent.


Applicants respectfully assert that the Chirashnya patent does not anticipate the invention as claimed. Chirashnya does not teach the claimed steps or means for defining an original link speed factor for each link, performing at least one runtime measurement for each link, calculating a runtime link speed factor for each link, and comparing the calculated runtime link speed factor to the original link speed factor. Applicants reiterate that "link speed" is carefully chosen language which does not encompass all attributes for characterizing a network. The claimed invention expressly defines, measures, calculates and compares link speeds. In contrast, Chirashnya sends packets through a switch adapter and then counts or evaluates the integrity of packets at the destination. Since Chirashnya does not define an original link speed, does not perform runtime measurements of link speed indicators and calculate runtime link speeds based on those measurements, and does not compare original to runtime link speeds, it cannot be maintained that Chirashnya anticipates the language of independent Claims 1, 17 and 28, or the claims which depend therefrom and add limitations thereto (Claims 2-8, 18-22), or those claims which recite parallel limitations (Claim 15).

Similarly with regard to Claims 11-16, 25-27, 30, and the other claims, Claims 4-6, 8, and 20, which depend from Claims 1 and 17 and which also recite application-based response to detected slow links, Applicants respectfully assert that the Chirashnya patent does not anticipate the claim language. Chirashnya does not teach application-based response to faulty switch adapters. At best, in response to identification of a faulty switch adapter, packets are re-routed under the Chirashnya patent. Clearly, therefore, Chirashnya does not anticipate the invention as set forth in claims 4-6, 8, 11-16, 20, 25-27, and 30.

Based on the foregoing amendments and remarks, Applicants respectfully request entry of the amendments, reconsideration of the rejections, and issuance of the claims.

Respectfully submitted,

Lorin Ullmann, et al

By: 
Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910